Bayes-factor of the ATLAS diphoton excess

Andrew Fowlie
ARC Centre of Excellence for Particle Physics at the Tera-scale,
School of Physics and Astronomy, Monash University, Melbourne, Victoria 3800 Australia
(Dated: December 7, 2016)

We present a calculation of Bayes-factors for the digamma resonance (F) versus the SM in light of ATLAS 8 TeV 20.3/fb, 13 TeV 3.2/fb and 13 TeV 15.4/fb data, sidestepping any difficulties in interpreting significances in frequentist statistics. We matched, wherever possible, parameterisations in the ATLAS analysis. We calculated that the plausibility of the F versus the Standard Model increased by about eight in light of the 8 TeV 20.3/fb and 13 TeV 3.2/fb ATLAS data, somewhat justifying interest in F models. All told, however, in light of 15.4/fb data, the F was disfavoured by about 0.7.

arXiv:1607.06608v2 [hep-ph] 6 Dec 2016

I. INTRODUCTION

The statistical anomalies at about 750 GeV in ATLAS [1, 2] and CMS [3, 4] searches for a diphoton resonance (denoted in this text as F) at $\sqrt{s} = 13\,\text{TeV}$ with about 3/fb caused considerable activity (see e.g., Ref. [5–7]). The experiments reported local significances, which incorporate a look-elsewhere effect (LEE, see e.g., Ref. [8, 9]) in the production cross section of the F, of 3.9σ and 3.4σ, respectively, and global significances, which incorporate a LEE in the production cross section, mass and width of the F, of 2.1σ and 1.6σ, respectively. There was concern, however, that an overall LEE, accounting for the numerous hypothesis tests of the SM at the LHC, cannot be incorporated, and that the plausibility of the F was difficult to gauge.

Whilst ultimately the F was disfavoured by searches with about 15/fb [10, 11], we directly calculate the relative plausibility of the SM versus the SM plus F in light of ATLAS data available during the excitement, matching, wherever possible, parameter ranges and parameterisations in the frequentist analyses. The relative plausibility sidesteps technicalities about the LEE and the frequentist formalism required to interpret significances. We calculate the Bayes-factor (see e.g., Ref. [12]) in light of ATLAS data,

$$\frac{p(\text{ATLAS data} \mid \text{SM} + F)}{p(\text{ATLAS data} \mid \text{SM})} = \frac{\frac{p(\text{SM} + F \mid \text{ATLAS data})}{p(\text{SM} \mid \text{ATLAS data})}}{\frac{p(\text{SM} + F)}{p(\text{SM})}}. \tag{1}$$

Our main result is that we find that, at its peak, the Bayes-factor was about 7.7 in favour of the F. In other words, in light of the ATLAS 13 TeV 3.2/fb and 8 TeV 20.3/fb diphoton searches, the relative plausibility of the F versus the SM alone increased by about eight. This was "substantial" on the Jeffreys' scale [13], lying between "not worth more than a bare mention" and "strong evidence." For completeness, we calculated that this preference was reversed by the ATLAS 13 TeV 15.4/fb search [11], resulting in a Bayes-factor of about 0.7. Nevertheless, the interest in F models in the interim was, to some degree, supported by Bayesian and frequentist analyses. Unfortunately, CMS performed searches in numerous event categories, resulting in a proliferation of background nuisance parameters and making replication difficult without cutting corners or considerable computing power.

II. CALCULATION

The background shape was characterised by a monotonically decreasing function with two free parameters,

$$p_b(m_{\gamma\gamma}) \propto \left[1 - \left(\frac{m_{\gamma\gamma}}{\sqrt{s}}\right)^{\frac{1}{3}}\right]^b \left(\frac{m_{\gamma\gamma}}{\sqrt{s}}\right)^a. \tag{2}$$

The $\sqrt{s} = 8\,\text{TeV}$ and $\sqrt{s} = 13\,\text{TeV}$ backgrounds were described by separate choices of a, b and normalisation, n_b.

ATLAS modelled the experimental resolution of the signal shape with a double-sided crystal ball (DSCB) function [1, 2]. In their combined analysis [1], ATLAS accounted for a substantial width by promoting DSCB parameters to functions of the mass and width of the F. Because the details of this treatment were not published, we picked a simpler ansatz for the signal shape and experimental resolution. The F signal was described by a Breit-Wigner or a Gaussian with a width equal to the ATLAS diphoton resolution if the width, Γ_F, was narrower than the ATLAS diphoton resolution, σ,

$$p_s(m_{\gamma\gamma}) \propto \begin{cases} \dfrac{1}{(m_{\gamma\gamma}^2 - m_F^2)^2 + \Gamma_F^2 m_F^2} & \Gamma_F > \sigma \\ e^{-\frac{(m_{\gamma\gamma} - m_F)^2}{2\sigma}} & \Gamma_F \leq \sigma \end{cases}, \tag{3}$$

and normalisation factors, n_s, at 8 TeV and 13 TeV. The ATLAS diphoton resolution was modelled by a linear function of F mass,

$$\sigma \approx 6 \cdot 10^{-3} m_F + 0.8\,\text{GeV}, \tag{4}$$

motivated by information in Ref. [1] that it changes from 2 GeV to 13 GeV between masses of 200 GeV and 2 TeV. We described the normalisation factors by an expected number of signal events in the 13 TeV 3.2/fb search and

SM		$\digamma$	
Parameter	Prior	Parameter	Prior
a_{13}, a_8	Flat, $[-25, 25]$	$m_{\digamma}/1\,\text{TeV}$	Log, $[0.2, 2]$
b_{13}, b_8	Flat, $[-25, 25]$	α	Log, $\left[5 \cdot 10^{-6}, 0.1\right]$
n_{b13}	Log, $\left[5 \cdot 10^3, 10^4\right]$	$\sigma_{13/8}$	Flat, $[2.5, 5]$
n_{b8}	Log, $\left[2 \cdot 10^4, 3 \cdot 10^4\right]$	n_{s13}	Log, $[5, 200]$

TABLE I: Priors for the SM ansatz and $\digamma$ resonance. Subscript numbers refer to 8 TeV and 13 TeV 3.2/fb, e.g., n_{b13} refers to the number of expected background events at 13 TeV with 3.2/fb.

	Evidence		
Data	SM	SM + $\digamma$	Bayes-factor
ATLAS 8 TeV 20.3/fb	$2.4 \cdot 10^{-64}$	$1.7 \cdot 10^{-64}$	0.71
ATLAS 13 TeV 3.2/fb	$6.8 \cdot 10^{-64}$	$3.1 \cdot 10^{-63}$	4.6
ATLAS 13 TeV 15.4/fb	$2.8 \cdot 10^{-87}$	$7.2 \cdot 10^{-88}$	0.26
8 TeV + 13 TeV 3.2/fb	$1.7 \cdot 10^{-127}$	$1.3 \cdot 10^{-126}$	7.7
8 TeV + 13 TeV 15.4/fb	$6.8 \cdot 10^{-151}$	$5.0 \cdot 10^{-151}$	0.73

TABLE II: Evidences for the SM ansatz and $\digamma$ resonance. Bayes-factors of greater than one indicate that the $\digamma$ is favoured.

scaling factors, reflecting the decreased cross section at $\sqrt{s} = 8\,\text{TeV}$ and different integrated luminosities. Thus the SM ansatz and $\digamma$ ansatz were described by six and four parameters, respectively.

We scraped the ATLAS bin counts, $\{n\}$, and bin edges from Ref. [2, 11]. To convert the distributions into expected numbers of events per bin, $\{\lambda\}$, we used analytic integration. Our likelihood function was simply a product of Poisson distributions,

$$\mathcal{L} = \prod_i \frac{\lambda_i^{n_i} e^{-\lambda_i}}{n_i!}, \tag{5}$$

for the expected and observed number of events in each bin in the 8 TeV and 13 TeV searches. In total, our calculations required about 6 million calls of our likelihood function.

The priors for the parameters were the final ingredients. We picked logarithmic priors for $m_{\digamma}$ and $\alpha \equiv \Gamma_{\digamma}/m_{\digamma}$ that matched the ranges in the ATLAS search: a mass between 200 GeV and 2000 GeV, and α between $5 \cdot 10^{-6}$ and 0.1. The latter range spans the narrow-width approximation (NWA) to substantial widths. We picked a logarithmic prior for the expected number of $\digamma$ events in the 13 TeV 3.2/fb search, n_{s13}, between 5 and 200 events, which reflected the $\digamma$ models anticipated in the experimental search. The number of expected events in the 8 TeV search was modelled by a scaling factor between the 13 TeV and 8 TeV production cross sections, $\sigma_{13/8}$. We picked a linear prior, i.e., $p(\sigma_{13/8}) = \text{const.}$, between 2.5 (corresponding to a light quark initial state) and 5 (corresponding to gluon fusion). We, of course, included a factor reflecting the decreased integrated luminosity. Since the models were composite (that is, we considered the SM and SM plus $\digamma$), we anticipated limited sensitivity to the priors for the SM background ansatz. We list all priors in Table I and discuss them further in Sec. II A.

We supplied the likelihood and priors to `MultiNest`[1] [14–16], which performs numerical integration via the nested sampling algorithm [17, 18], returning a Bayesian evidence, e.g.,

$$p(\text{ATLAS data} \mid \text{SM}) = \int \mathcal{L}(\boldsymbol{x}) \cdot p(\boldsymbol{x} \mid \text{SM})\, d\boldsymbol{x}, \tag{6}$$

where the factors in the integrand are the likelihood and prior, respectively, as discussed above, and $\boldsymbol{x}$ denotes the SM ansatz parameter set (see Appendix A) for a complete expression). Finally, we calculated Bayes-factors, which are ratios of Bayesian evidences. For a clear picture of the changing relative plausibility of the $\digamma$ versus the SM, we calculated Bayes-factors for 8 TeV, 13 TeV 3.2/fb and 13 TeV 15.4/fb separately and combined. We validated our likelihood function and scanning by checking that we approximately reproduced the best-fit $\digamma$ properties and significances reported by ATLAS; see Fig. 1 and Fig. 2. We found local significances of 3.9σ and 2.1σ at 13 TeV 3.2/fb and 8 TeV, respectively, and a combined local significance of 4.1σ.[2] We found best-fits for $(m_{\digamma}, \alpha)$ at about $(745\,\text{GeV}, 0.06)$, $(717\,\text{GeV}, 0.08)$ and $(739\,\text{GeV}, 0.09)$ at 13 TeV 3.2/fb, 8 TeV and combined, respectively. We, furthermore, found 1.4σ tension between the preferred ratio of cross sections in 13 TeV 3.2/fb and 8 TeV data with mass and width fixed to their 13 TeV 3.2/fb best-fits. We found no indication that our signal model in Eq. (3) was an inadequate or poor approximation to the unknown DSCB function.

We calculated that the $\sqrt{s} = 8\,\text{TeV}$ data slightly disfavours the $\digamma$ by a Bayes-factor of about 0.7. At $\sqrt{s} = 13\,\text{TeV}$ with 3.2/fb, the story changes. The $\digamma$ is favoured by about five with the region around $m_{\digamma} \approx 750\,\text{GeV}$ dominating, as expected. We calculated that the 8 TeV and 13 TeV 3.2/fb ATLAS data combined favoured the $\digamma$ by about 7.7. This is greater than a naive multiplication of Bayes-factors; the Bayes-factors cannot be combined by multiplication, as the evidences are dependent. All told, however, a combination of 8 TeV and 13 TeV with 15.4/fb disfavours the $\digamma$ by about 0.7. Whilst this is evidence against the $\digamma$, it is "not worth more than a

[1] We picked an evidence tolerance of 0.01 and 1000 live points per dimension.

[2] We assumed that the log-likelihood ratio was $\frac{1}{2}\chi_1^2$-distributed (see e.g., Ref. [19]).

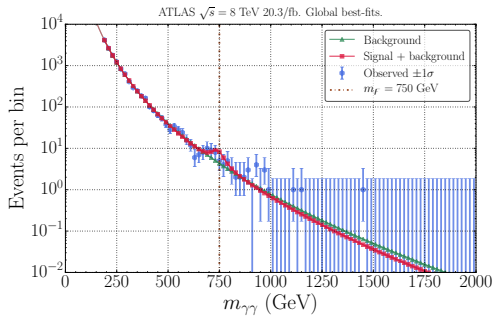

(a) Combined best-fit and spectrum at $\sqrt{s} = 8\,\text{TeV}$.

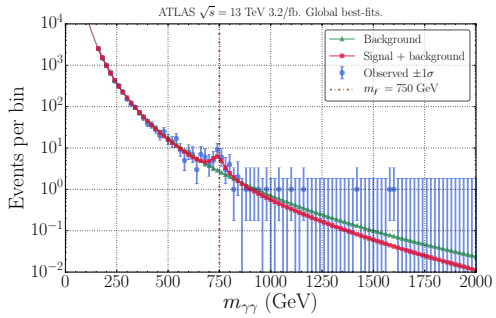

(b) Combined best-fit and spectrum at $\sqrt{s} = 13\,\text{TeV}$ 3.2/fb.

FIG. 1: Diphoton spectrum at (a) $\sqrt{s} = 8\,\text{TeV}$ and (b) $\sqrt{s} = 13\,\text{TeV}$ 3.2/fb, along with best-fit background and signals.

bare mention" on the Jeffreys' scale. The evidences and Bayes-factors are summarised in Table II.

A. Prior sensitivity of the Bayes-factor

There is, of course, that thorny issue that a Bayes-factor is a functional of our priors for the models' parameters. This is most dangerous in cases in which an evidence is sensitive to the interval of a prior for a parameter. Consider a model with a single parameter λ with a linear prior in the interval $\Delta\lambda$, and suppose that the likelihood is reasonable only in the interval σ_λ. The evidence would be approximately proportional to $\sigma_\lambda/\Delta\lambda$. For robust inference, it appears that we require prior information indicating a plausible interval. In fact, this is not strictly necessary. Such factors could cancel in a Bayes-factor, even for improper priors, i.e., $\Delta\lambda \to \infty$. Furthermore, pragmatically, one could calculate an evidence for a sub-model with λ in a region of interest (e.g., a region searched by an experiment), $\delta\lambda$, learning whether such a sub-model was confirmed by data. This, in effect, shifts any $1/\Delta\lambda$ factors from an evidence into a prior for a sub-model, i.e.,

$$p(\text{sub-model}) = p(\text{model}) \cdot p(\lambda \in \delta\lambda \mid \text{model}) \propto \delta\lambda/\Delta\lambda. \tag{7}$$

The resulting Bayes-factor indicates the relative change in plausibility of the sub-model.

There may, furthermore, be cases in which our prior information cannot uniquely determine a distribution for a parameter upon an interval, i.e., several choices of distribution may appear consistent with our prior information. This is problematic; the different choices may lead to different Bayes-factors. However, remarkably, in many cases particular ignorance about a parameter uniquely determines a distribution (see e.g., Ref. [21]). If we are ignorant, e.g., of the scale of λ, our prior should be invariant under $\lambda \to A\lambda$, leading to a logarithmic prior, $p(\log\lambda) = \text{const.}$[3]

Let us consider these issues in detail for all our priors. To quantify sensitivity, we recalculated Bayes-factors for the 8 TeV + 13 TeV 3.2/fb data:

- Fortunately, because the SM background ansatz was common to each model, any such factors originating from SM background parameters would vanish in a ratio of evidences, i.e., a Bayes-factor. Similarly, we anticipated limited sensitivity to the shape of the priors for the SM background ansatz parameters.

- Whilst this was not the case for the F mass and width, narrower intervals than those in Table I would imply that we were in possession of prior information that precluded resonances in regions that were searched by ATLAS. Wider intervals would merely damage the plausibility of the F model by diluting its evidence, though extreme $\alpha \equiv m_F/\Gamma_F$ could be implausible from the perspective of QFT. Our intervals matched the resonance masses and widths searched for by ATLAS, i.e., we considered the change in plausibility of a F resonance searched for by ATLAS. We picked logarithmic priors for the mass and width of the F; anything else would imply prior information favouring particular scales in the intervals in Table I.

  Nevertheless, with linear priors (which imply that prior information favoured the highest scales permissible) for the F mass and $\alpha \equiv m_F/\Gamma_F$, we find a Bayes-factor of 39.7 in favour of the F.[4] This is about 5 times greater than previously; dominantly because a linear prior favoured the best-fitting $\alpha \approx 0.09$ by about 10 relative to a logarithmic prior. From a theoretical perspective, however, a substantial width was, if anything, implausible relative to a narrow width; if we had any reliable

[3] Whilst this distribution is improper, in practice we may be in possession of prior information that λ cannot be arbitrarily great or small, though we are ignorant of its scale within that interval, leading to a proper distribution.

[4] This Bayes-factors was found by nested sampling. All further Bayes-factors were found by re-weighting an existing chain.

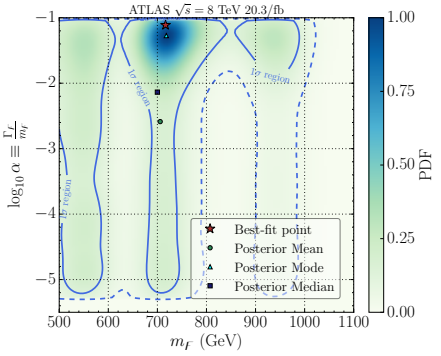

(a) ATLAS $\sqrt{s} = 8$ TeV 20.3/fb data.

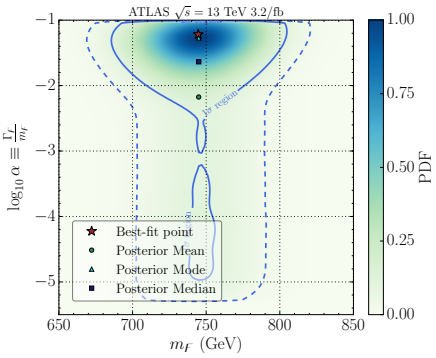

(b) ATLAS $\sqrt{s} = 13$ TeV 3.2/fb data.

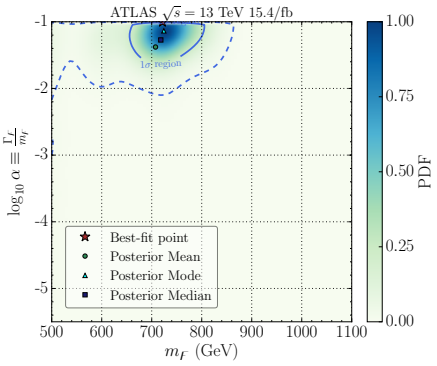

(d) ATLAS $\sqrt{s} = 13$ TeV 15.4/fb data.

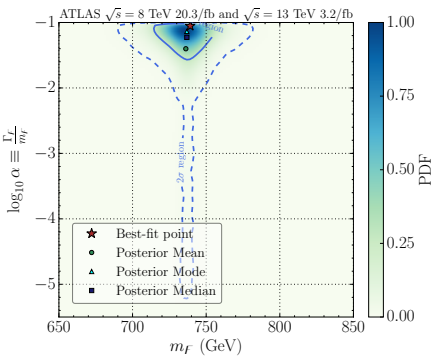

(c) ATLAS $\sqrt{s} = 8$ TeV 20.3/fb and 13 TeV 3.2/fb data.

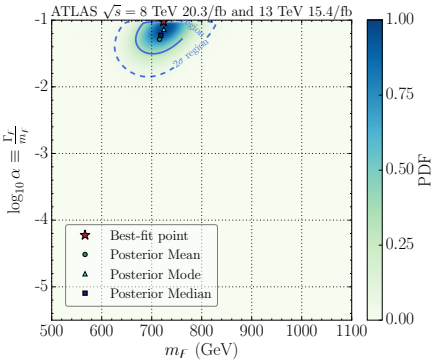

(e) ATLAS $\sqrt{s} = 8$ TeV 20.3/fb and 13 TeV 15.4/fb data.

FIG. 2: The posterior pdf (see Ref. [20]) for the mass and width of F. Left-panels: (a) ATLAS $\sqrt{s} = 8$ TeV data, (b) 13 TeV 3.2/fb data and (c) combined. Right panels: (d) 13 TeV 15.4/fb data and (e) 8 TeV and 13 TeV 15.4/fb combined.

prior information about the scale of the F width, it was that should be narrow.

- We picked a linear prior for the ratio of the 13 TeV and 8 TeV cross-sections. The interval spanned 2.5, corresponding to a light-quark initial state, to 5, corresponding to gluon fusion. This prior was motivated by knowledge about plausible production mechanisms; we were not ignorant of its scale and a light-quark initial state was a priori as plausible as gluon fusion. Nevertheless, we found that a logarithmic prior, which implies that prior information favoured a light-quark initial state, reduced the Bayes-factor by a factor of about 0.9.

- We picked a logarithmic prior between 5 and 200 for the number of F signal events in the 13 TeV 3.2/fb search. Whilst we were ignorant of scale within an interval, an extreme number of signal events was implausible from the perspective of QFT (an extreme cross section) and experiments (no other evidence for a resonance with an extreme cross section), resulting in an upper limit. Reducing the maximum number of events to 50 increased the

Bayes-factor by a factor of about 1.6. Reducing the minimum number of events would, asymptotically, result in an SM background-like model and thus a Bayes-factor of 1.

We stress that our interpretation is that our Bayes-factor is the change in relative plausibility of a F resonance searched for by ATLAS, i.e., with a mass in the interval 200 GeV to 2 TeV and $\alpha \equiv m_F/\Gamma_F$ in the interval $5\cdot10^{-6}$ to 0.1 (see Table I), and that the priors chosen reflected knowledge about cross sections and widths in QFT, and, in some cases, ignorance of scale or location. The Bayes-factor increased by a factor of about 5 in the worst-case — linear priors of the F mass and width; however, it was a somewhat dishonest choice, since it implied that prior information favoured an appreciable F mass and width. In other words, the prior information was sufficient to insure a weak dependence on choices of prior.

B. Posterior distributions of F properties

The posterior pdfs for the F properties were by-products of our calculations. Considering the combined 8 TeV and 13 TeV 3.2/fb data, for the mass, the posterior mean, median and mode were about $m_F \approx 737$ GeV, with a symmetric two-tailed 68% credible region spanning 724 GeV to 747 GeV. For the width, the posterior mean, median and mode in $\log\alpha$ differed, but spanned about $\alpha \approx 0.05$ to 0.08, with a one-tailed 1σ (2σ) lower bound at 0.05 (0.004). Finally, for the ratio of cross sections, the posterior pdf favoured $\sigma_{13/8} \approx 5$, corresponding to gluon fusion, but smaller ratios were permitted with a one-tailed 1σ (2σ) lower bound at 3.8 (2.8). In other words, the posterior pdf favoured a mass of about 740 GeV, a large width and production by gluon fusion, as expected.

We show posterior pdf on the (m_F, Γ_F) plane in Fig. 2 for 8 TeV, 13 TeV 3.2/fb and 13 TeV 15.4/fb data separately and combined. The credible regions at $m_F \lesssim 500$ GeV (not shown) were vulnerable to digitisation errors in the data scraped from the low-mass region in Ref. [2, 11]. Fortunately, that region ultimately contributed little to the evidences or our conclusions. The three prongs in the pdf at 8 TeV in Fig. 2a originated from deficits and excesses that surrounded the excess near 750 GeV in Fig. 1a. The pdf at 13 TeV 3.2/fb in Fig. 2b exhibited a single prong around 750 GeV that narrowed once data was combined in Fig. 2e.

III. CONCLUSIONS

Statistical anomalies near 750 GeV in searches for diphoton resonances at the LHC resulted in a frenzy of model building. In the official analyses, the data was investigated with frequentist techniques. To sidestep issues regarding the interpretation of significances, with Bayesian statistics, we calculated the change in plausibility of the F resonance versus the SM in light of the ATLAS 8 TeV 20.3/fb, 13 TeV 3.2/fb and 13 TeV 15.4/fb diphoton searches. There was limited freedom in the choice of priors for the F: we matched, where possible, the ranges of width and mass searched by ATLAS. Since the models were composite, we expected limited sensitivity to the priors for the SM background ansatz and found weak dependence on our choice of priors for the F signal. Ideally, we would have combined ATLAS and CMS data, though the numerous categories in the latter make a combination challenging. Our F was a toy-model described by a simple Breit-Wigner with mass, width and cross section potentially within reach of the ATLAS search; we would find different Bayes-factors for a theory that precisely predicted the F properties (though know of no such theory).

We calculated that the relative plausibility of the F increased by about 7.7 in light of the ATLAS data available at the height of the excitement. This should be contrasted with conclusions from frequentist analysis, e.g., the probability of obtaining a test statistic as extreme as that observed in the 13 TeV 3.2/fb search were the SM correct was about 0.02 (2.1σ). This Bayes-factor was unimpressive, especially considering that e.g., SM precision measurements could quash that preference and that a width of $\Gamma_F \approx 0.06 m_F$ was somewhat unexpected, a fact not reflected by our priors. On the other hand, a combination with CMS data could have increased the Bayes-factor, though there may have been tension in the preferred width. Considering all data, including 13 TeV 15.4/fb, the F was disfavoured by about 0.7. As well as aiding our understanding of the 750 GeV anomaly, we hope our calculations serve as a proof of principle for Bayesian model comparison and parameter inference in future experimental searches.

Appendix A: Evidence calculation

An expression for an evidence was written schematically in Eq. (6). As an example, we now write in detail the evidence of the SM + F model in light of the ATLAS $\sqrt{s} = 13$ TeV 3.2/fb data. All other evidence integrals were performed in a similar manner. We begin from the usual expression for the evidence (see e.g., Ref. [13]),

$$\begin{aligned}\mathcal{Z} &\equiv p\big(\text{ATLAS } \sqrt{s} = 13\,\text{TeV } 3.2/\text{fb} \mid \text{SM} + F\big) \\ &= \int \mathcal{L}(\boldsymbol{y}) \cdot p(\boldsymbol{y} \mid \text{SM} + F)\, d\boldsymbol{y},\end{aligned} \tag{A1}$$

where $\boldsymbol{y}$ denotes $\{a_{13}, b_{13}, n_{b13}, n_{s13}, m_F, \alpha\}$, i.e., the SM + F parameters, the likelihood function, $\mathcal{L}(\boldsymbol{y})$, was a product of Poissons (Eq. (5)) and the priors were independent. Explicitly,

$$\cdots = \int \prod_i \frac{\lambda_i(\boldsymbol{y})^{n_i} e^{-\lambda_i(\boldsymbol{y})}}{n_i!} \cdot \prod_j p(y_j \mid \text{SM} + F)\, dy_j, \tag{A2}$$

where i indexes diphoton mass bins in the ATLAS $\sqrt{s} = 13\,\text{TeV}$ 3.2/fb search, such that n_i is the number of observed events in bin-i, and j indexes the SM + F parameters. The expected number of events per bin, λ_i, was a function of the SM + F parameters,

$$\lambda_i(\boldsymbol{y}) = \int_{\text{bin-}i} n_{b13} \cdot p_b(m_{\gamma\gamma}; \boldsymbol{y}) + n_{s13} \cdot p_s(m_{\gamma\gamma}; \boldsymbol{y})\, dm_{\gamma\gamma}, \tag{A3}$$

where the p_b and p_s are the background and signal diphoton distributions in Eq. (2) and Eq. (3), respectively. The product of priors was

$$\prod_j p(y_j \mid \text{SM} + F) \propto \frac{1}{n_{s13} n_{b13}} \frac{1}{m_F\, \alpha} \tag{A4}$$

inside the intervals in Table I and zero elsewhere. The reciprocal factors were for logarithmic priors. The priors were, of course, normalised such that

$$\int \prod_j p(y_j \mid \text{SM} + F)\, dy_j = 1. \tag{A5}$$

The integration in Eq. (A3) was performed analytically; all other integration was performed numerically with the nested sampling Monte-Carlo algorithm [17, 18].

[1] ATLAS, ATLAS-CONF-2015-081, (2015).
[2] ATLAS, M. Aaboud *et al.*, JHEP **09**, 001 (2016), arXiv:1606.03833 [hep-ex].
[3] CMS, CMS-PAS-EXO-15-004, (2015).
[4] CMS, V. Khachatryan *et al.*, Phys. Rev. Lett. **117**, 051802 (2016), arXiv:1606.04093 [hep-ex].
[5] J. Ellis, S. A. R. Ellis, J. Quevillon, V. Sanz, and T. You, JHEP **03**, 176 (2016), arXiv:1512.05327 [hep-ph].
[6] R. Franceschini *et al.*, JHEP **03**, 144 (2016), arXiv:1512.04933 [hep-ph].
[7] A. Strumia, (2016), arXiv:1605.09401 [hep-ph].
[8] L. Lyons, (2013), arXiv:1310.1284 [physics.data-an].
[9] L. Lyons, Ann. Appl. Stat. **2**, 887 (2008), arXiv:0811.1663 [stat.AP].
[10] CMS, CMS-PAS-EXO-16-027, (2016).
[11] ATLAS, ATLAS-CONF-2016-059, (2016).
[12] P. Gregory, *Bayesian Logical Data Analysis for the Physical Sciences* (Cambridge University Press, 2005).
[13] H. Jeffreys, *The Theory of Probability* (Oxford University Press, 1939).
[14] F. Feroz and M. P. Hobson, Mon. Not. Roy. Astron. Soc. **384**, 449 (2008), arXiv:0704.3704 [astro-ph].
[15] F. Feroz, M. P. Hobson, and M. Bridges, Mon. Not. Roy. Astron. Soc. **398**, 1601 (2009), arXiv:0809.3437 [astro-ph].
[16] F. Feroz, M. P. Hobson, E. Cameron, and A. N. Pettitt, (2013), arXiv:1306.2144 [astro-ph.IM].
[17] J. Skilling, Bayesian Anal. **1**, 833 (2006).
[18] J. Skilling, AIP Conference Proceedings **735**, 395 (2004).
[19] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, Eur. Phys. J. **C71**, 1554 (2011), arXiv:1007.1727 [physics.data-an], [Erratum: Eur. Phys. J.C73,2501(2013)].
[20] A. Fowlie and M. H. Bardsley, (2016), arXiv:1603.00555 [physics.data-an].
[21] E. T. Jaynes, *Probability Theory: The Logic of Science* (Cambridge University Press, 2003).